

PCTWORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 9/46	A1	(11) International Publication Number: WO 97/29424 (43) International Publication Date: 14 August 1997 (14.08.97)
<p>(21) International Application Number: PCT/US97/01474</p> <p>(22) International Filing Date: 5 February 1997 (05.02.97)</p> <p>(30) Priority Data: 08/598,800 9 February 1996 (09.02.96) US</p> <p>(71) Applicant: CITRIX SYSTEMS, INC. [US/US]; 6400 N.W. 6th Way, Fort Lauderdale, FL 33309 (US).</p> <p>(72) Inventor: PEDERSEN, Bradley, J.; 1335 N.W. 112 Terrace, Coral Springs, FL 33071 (US).</p> <p>(74) Agent: LANZA, John, D.; Testa, Hurwitz & Thibault, L.L.P., High Street Tower, 125 High Street, Boston, MA 02110 (US).</p>		<p>(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ARIPO patent (KE, LS, MW, SD, SZ, UG), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).</p> <p>Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p>
(54) Title: METHOD AND APPARATUS FOR CONNECTING A CLIENT NODE TO A SERVER NODE BASED ON LOAD LEVELS		
<p>(57) Abstract</p> <p>A network including at least two server nodes, a master network information server node and a client node in communication with each other. The master network information server node contains a list of network addresses and respective load information corresponding to each of the server nodes. In one embodiment the client node is directed to use a server node on which to execute an application based on the load information stored on the master network information server node. A method is described which includes the steps of requesting an available server from a master network information server node and connecting the client node to an available server node based on the load levels.</p>		

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LJ	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

METHOD AND APPARATUS FOR CONNECTING A CLIENT NODE TO A SERVER NODE BASED ON LOAD LEVELS

Field of the Invention

The invention relates generally to a method and apparatus for computer to computer
5 communication and more specifically for communication between a client node and a server node
over a communication network.

Background of the Invention

Computer communication networks typically include one or more computer nodes termed
servers and one or more computer nodes termed clients. A server provides a service to a client
10 upon receiving a request from the client. An example of a typical service provided by a server is
that provided by a file server.

A file server is a computer node having data storage capabilities. Data files which are to
be accessed by users on various client nodes are stored on the file server. When a client node
needs a file, a request is transmitted to the file server node which obtains the file from its local
15 storage and transmits the file to the requesting client node. When the client node has completed
its use of the file, the client node returns the updated file to the file server for storage. By using a
file server rather than keeping files on local storage on the client node, the amount of local storage
is reduced, the ability to backup or copy the file periodically is improved, and access to the file by
multiple users is improved.

20 Many other types of servers are known. For example, print servers control the printing of
documents for client nodes. By using a print server, each computer need not have a printer,
thereby reducing the number of printers needed to support the number of computers on the
network. Application servers provide application programs for clients, thereby avoiding the need

- 2 -

of maintaining an application on each client. Application execution servers not only provide applications for the client nodes, but also execute the applications and thereby act as "compute engines" for the client nodes.

One issue that arises is that of apportioning the client requests to each server node which is capable of answering the requests. For example, if two application execution server nodes are both capable of executing an application for a client node, but one application execution server node is presently executing the application for three other clients and the other application server node is not executing the application for any client, then what is desired is a means for sending any new requests for the execution of the application to the under-utilized server. This balancing of the work load among the various network servers is referred to as load leveling.

One prior art system which is capable of performing a rudimentary form of load leveling is the VAXcluster™ computer system of the Digital Equipment Corporation (Maynard, Massachusetts). A VAXcluster™ is a group of VAX™ computers connected together by a communication link. A user on a local area network connected to the VAXcluster™ can access different computers of the VAXcluster™ according to the load level of the various nodes. However, this form of load leveling is more akin to a user selecting a different client node than it is selecting a server node in a present day network.

What is desired then is a method for permitting a client node to select the least loaded server node without user intervention. The present invention permits such functionality.

Summary of the Invention

The invention relates to a network including at least one server node, a network information server node and a client node in communication with each other. The network information server node includes a list of network addresses and load information corresponding to each the server nodes. In one embodiment the client node is directed to a server node on

- 3 -

which to execute an application based on the list of available servers and corresponding load levels as determined by the network information server node. In one embodiment the network information server node is a separate node from the other server nodes. In another embodiment the master network information server node is also an application execution server node. In yet another embodiment the network comprises at least two network information server nodes, one of the network information server nodes being elected the master network information server node. In another embodiment the application execution server nodes are grouped into clusters.

The invention also relates to a method for connecting a client node to a server node. In one embodiment, the method includes the steps of requesting an available service and being directed to a server based upon the load levels as determined by a network information server node.

Brief Description of the Drawings

This invention is pointed out with particularity in the appended claims. The above and further advantages of this invention may be better understood by referring to the following description taken in conjunction with the accompanying drawings, in which:

Fig. 1 is a diagram of an embodiment of a local network on which the invention is practiced;

Fig. 2 is a diagram of an embodiment of a local network on which the invention is practiced, with the application execution servers grouped into clusters;

Fig. 3 is a diagram of an embodiment of a wide area network on which the invention is practiced;

Fig. 4 is a flowchart representation of an embodiment of a process for electing a master network information server node;

- 4 -

Fig. 5 is a flowchart representation of an embodiment of a process by which servers update the master network information server node; and

Fig. 6 is a flowchart representation of a request by a client to an embodiment for making a master network information server node and a server node.

Detailed Description of the Invention

In brief overview, and referring to Fig. 1, a typical network 10 includes a client node 18, two server nodes 26, 26', and a master network information server node 34 connected together by a communications link 42. The embodiment shown in Fig. 1 depicts the communications link 42 as a local area network or LAN, but any communication topology may be used. For the purpose of simplification, the two server nodes 26, 26' are both capable of servicing a request from the client node 18. Also for the purpose of simplification the master network information server node 34 is assumed to be a distinct server node, but in actuality the master network information server node 34 may be an application execution server node 26, 26'. On a given LAN several nodes may be capable of acting as a network information server node, but only one of such nodes is designated the master network information server node and it is to this node that client requests for server information are directed.

The master network information server node 34 maintains a table of addresses for the application execution server nodes 26, 26'. In addition, the master network information server node 34 receives messages from the application execution server nodes 26, 26' indicating their level of activity. The level of activity of the application execution server nodes 26, 26' is maintained in a table along with the address of each of the application execution server nodes 26, 26'.

When a client 18 wishes to have an application executed on an application execution server node 26, 26', the client node 18 sends a request to the master network information server node 34. The master network information server node 34 accesses the table of server addresses and returns a message containing the address of the application execution server 26, 26' with the least load. The client node 18 then requests service from this least loaded application execution server node 26, 26'.

- 6 -

Referring to Fig. 2, the application execution servers 26, 26', 26'' and the master network information servers 34, 34' may be grouped into clusters 50, 50'. A cluster 50, 50' is defined as a group of application execution servers 26, 26', 26'' with similar capabilities. For example, each application execution server 26, 26', 26'' in a cluster 50, 50' is capable of executing a given application. The cluster 50, 50' is named and is accessed by a client 18, by using the cluster name. The cluster name can be the name of an application common to all members of the cluster 50, 50'. A given application execution server 26, 26', 26'' can belong to more than one cluster 50, 50'. Thus, if application execution server 26 is capable of executing applications A and B, server 26' is capable of executing applications B and C and server 26'' is capable of executing applications A and C, then cluster A would include servers 26 and 26'', cluster B would include servers 26 and 26' and cluster C would include servers 26' and 26''. All client nodes 18, application execution server nodes 26, 26', 26'' and master network information nodes 34, 34' which communicate do so using the same network transport layer communication protocol.

Therefore, if the client node 18, the application server nodes 26, 26' and master network information server node 34 communicate among themselves, the client 18, the servers 26, 26' of the cluster 50 and the master network information server 34 would communicate using the same protocol, for example TCP-IP. Similarly if client 18', the server node 26'' and the master network information server node 34' of cluster 50' communicate using a different protocol, for example IPX, then only client nodes using IPX can communicate with the server 26'' and the master network information server node 34'. Thus there must be a master network information server node 34 for every network protocol used.

Referring to Fig. 3, the system of the invention can also be used in the context of a Wide Area Network (WAN). In the embodiment of the WAN shown in Fig. 3, two local area networks 10', 10'' are connected by a WAN communication link 58. Each LAN 10', 10'' may include a

- 7 -

client node 18, 18''' and/or one or more application execution server nodes 26, 26', 26'', 26'''.

In addition each LAN may include one or more gateway nodes 64, 64'. A gateway node is a node on one LAN which maintains the address of at least one node on a remote LAN.

When there is a gateway node 64, 64' on the LAN 10, 10', that gateway node 64, 64' is
5 also the master network information server node 34 for the LAN 10'', 10', respectively. The master network information server node 64 does not have any information about the application execution servers 26''' located on the remote network 10'. All the master network information server node 64 has information about is the corresponding master network information server node 64' on the remote network 10', and the master network information server node 64
10 communicates with its counterpart 64' on the remote network 10' to obtain information about the application execution servers 26''' located on the remote network 10'. Thus, a master network information server node 64 configured as a gateway node maintains the address of the corresponding remote master network information server node 64'. As such, the master network information node 64 configured as a gateway node to remote network 10' maintains the addresses
15 of the remote master network information server nodes 64'.

Referring again to Fig. 1, It is desirable to dynamically select a master network information server node 34 from the available nodes on the network. In this way, if the master network information server node 34 fails, a new master network information server node 34 may be selected as soon as the failure of the previous master network information server node 34 is
20 detected. In one embodiment a master network information server node 34 is selected by an election process among the servers 26. Note that if the master network information server node 34 changes because of an election, the gateway node will change to the new master network information server 34.

- 8 -

In one embodiment, any node (client 18, server 26, or network information server 24) may force an election at any time by broadcasting a request election datagram on the LAN 10. The election results are determined by a comparison of the set of election criteria which is transmitted within the request election datagram transmitted by the requesting node with the set of election
5 criteria maintained on each receiving node. That is, the first election criterion from the datagram of the requesting node is compared by the receiving node to the first criterion of the receiving node. The highest ranking of the two criteria being compared wins the comparison and the node with that criterion wins the election. If the two criteria tie, then the next criteria are sequentially compared until the tie is broken. If server node 26 receiving the request election datagram has a
10 higher election criteria than that received in the request election datagram, the node receiving the request election datagram issues its own request election datagram. If the receiving node has a lower election criteria than the criteria received in the request election datagram, the receiving node determines it is not the master network information server node and attempts to determine which node on the network is the master network information server node.

15 In one embodiment the criteria which determine the outcome of the election include:
whether or not the node is statically configured as a master network information server node;
whether the node has the higher master network information server software version number;
whether the node is an NT domain controller (The domain controller is a server which maintains
the master user database for the network. Because this controller has high availability it is a good
20 selection for the network information server.); whether the node is the longest running node; and
whether the node has a lexically lower network name. In one embodiment, the datagram
structure for the election request includes an unsigned shortword for the server version number,
an unsigned shortword in which the bits are flags which designate whether the node is statically

- 9 -

configured as a master network information server node, or is executing on a NT domain controller and an unsigned longword containing the amount of time the server has been running.

Periodically, the master network information server node 34 transmits a declare message to the other nodes on the LAN declaring itself to be the master network information server node

- 5 34. If another node believes itself to be a master network information server node 34 for the cluster 50, the other node will request an election. In this way erroneous master network information server nodes 34 of the same protocol are detected and removed. In addition an election will also be requested: by any node when that node reboots; by any server node to whom the master network information server node has failed to acknowledge an update message; or any
- 10 client node to whom the master network information server node has failed to respond to a request for information.

- In more detail and referring to Fig. 4, once any node broadcasts a request election datagram requesting an election (Step 20), the node receiving the request election datagram (Step 24) first compares its election criteria to the criteria in the request election datagram (Step 30) to
- 15 determine if the receiving node has higher criteria (Step 34). If the node receiving the datagram has lower election criteria (Step 38) than the criteria contained in the request election datagram, the node receiving the request election datagram drops out of the election process and awaits the results of the election (Step 38).

- If the node receiving the request election datagram has higher election criteria than that
- 20 contained in the request election datagram, then the node receiving the request election datagram broadcasts its own request election datagram containing the nodes own election criteria (Step 40). If in response to the transmission of the request election datagram by the second node, another node responds with a request election datagram with even higher election criteria, then the second node drops out of the election and the node with higher criteria broadcasts it's own request

- 10 -

election datagram. If no other node responds with higher election criteria, the node which has apparently won the election for master network information server node sends n more election requests, (in one embodiment three requests) (Step 56) and then if still no other node responds with higher election criteria, the node which has sent the n election requests is the new master
5 network information server node.

After the election has occurred and the new network information server 34 has been determined, all the server nodes 26, 26', 26'' send all of their configured gateway addresses to the new network information server node 34. In this way the new network information server node 34 becomes a gateway node.

10 Referring again to Fig. 1, once the master network information server node 34 is elected, the application execution server nodes 26, 26', 26'' send update datagrams to the master network information server node 34 providing information about each server node 26, 26', 26'' transmitting the update datagram. In one embodiment, the update datagram sent to the master network information server node 34 from a server node 26, 26', 26'' includes: the server name;
15 the network address; the cluster name; the network transport protocol; the total number of nodes configured with this transport; the number of ports available for connection with a client using this transport protocol; the total number of users permitted to be active at one time; number of available user slots; and server load level. Upon receipt of the update datagram, the master network information server node 34 returns an acknowledgment to the server node 26, 26', 26''
20 which transmitted the update datagram indicating that the update datagram was received. If the server node 26, 26', 26'' transmitting the update datagram does not receive an acknowledgment from the master network information server node 34, the transmitting server node 26, 26', 26'' assumes that the master network information server node 34 has failed and transmits an election request.

- 11 -

In more detail and referring to Fig. 5, a server node 26, 26', 26'', after the election of a master network information server node 34, waits a random period of time and then sends a datagram to the master network information server node 34 with its latest load information (Step 100). In one embodiment the delay is between four and six seconds. If the master network information server node 34 receives (Step 108) an update datagram from a server node 26, 26', 26'', then the master network information server node 34 replies to the transmitting node with an acknowledgment (Step 110) and forwards the data to any node configured as a gateway node. If the master network information server 34 fails to receive data from a server 26, 26', 26'' (Step 108), then the master network information server 34 discards the old data from the server node 26, 26', 26'' after a predetermined amount of time (Step 120).

If the server 26, 26', 26'' node does not receive an acknowledgment from the master network information server node 34 after the server 26, 26', 26'' has sent an update datagram (Step 128), the server 26, 26', 26'' retransmits the update datagram. The server 26, 26', 26'' will attempt n retransmits (in one embodiment three) before it assumes that the master network information server 34 has failed and then transmits an election request (Step 130). If the server node receives an acknowledgment, then it periodically updates the master network information server node, in one embodiment every 5 to 60 minutes (Step 140).

Once the master network information server node 34 has established a database of information relating to the server nodes 26, 26', a client 18 may request information about the servers 26, 26'. Referring to Fig. 6, the client node 18 first determines which node is the master network information server node 34. To do so, the client node 18 broadcasts a datagram to all server nodes asking which node is the master network information server node 34 for the transport protocol used by the client 18 (Step 300). Upon receipt by a network information server node 34 (Step 306), the network information server node 34 transmits the identity of the

- 12 -

master network information server node to the client node 18 by datagram (Step 310). The client node 18 then receives the identity datagram from the network information server nodes 34 (Step 316) and using the address provided requests information regarding the servers on the network regarding the application desired by the client 34(Step 320). The master network information server node 34 receives the request (Step 326) and transmits the address of the least loaded server node 26, 26' to the requesting client 18 (Step 330). The client 18 receiving the address of the least loaded server node 26, 26' (Step 336), transmits a request to the server node 26, 26' (Step 338), which starts the application (Step 340) and establishes a connection with the client 18 (Step 350).

Having described preferred embodiments of the invention, it will now become apparent to one of skill in the art that other embodiments incorporating the concepts may be used. It is felt, therefore, that these embodiments should not be limited to disclosed embodiments but rather should be limited only the spirit and scope of the following claims.

CLAIMS

What is claimed is:

- 1 1. A network information server node comprising:
 - 2 a memory comprising:
 - 3 a plurality of network addresses and
 - 4 a plurality of load information, each of said plurality of load information
 - 5 corresponding to a respective one of said plurality of network
 - 6 addresses; and
 - 7 a process receiving client requests and responding to said client requests with
 - 8 network addresses and corresponding load information.
- 1 2. A network information server node comprising:
 - 2 a memory comprising a data table, said data table storing:
 - 3 a plurality of network addresses and load information associated with each of
 - 4 said plurality of network addresses; and
 - 5 a process receiving client requests and responding to said client requests with
 - 6 network addresses and associated load information.
- 1 3. A network information server node comprising:
 - 2 a memory comprising a data table, said data table having a plurality of entries, each
 - 3 said entry corresponding to a respective network node of a plurality of
 - 4 network nodes, each of said entries including:
 - 5 a network address of a respective one of said plurality of network nodes
 - 6 and

- 14 -

load information associated with said respective one of said plurality of
network nodes; and

a process receiving client requests and responding to each of said client requests
with entries of said data table corresponding to network nodes which
satisfy said each of said client requests.

4. A network comprising:

at least two server nodes;

a first information server node in communication with said at least two server nodes, said

first information server node comprising a memory comprising a data table, said
data table having an entry for each of said at least two server nodes, each of said
entries including:

an address corresponding to a respective one of said at least two server nodes; and
load information corresponding to said respective one of said at least two server
nodes; and

a client node in communication with said first information server node and said at least
two server nodes, said client node executing an application on one of said at least
two server nodes in response to said load information obtained from said first
information server node.

5. The network of claim 4 further comprising a plurality of client nodes in communication
with said first information server node.

6. The network of claim 4 further comprising at least two information server nodes wherein
one of said at least two information server nodes is a master network information server node.

7. The network of claim 4 wherein said at least two server nodes are grouped into clusters.

- 15 -

1 8. The network of claim 6 wherein one of said at least two information server nodes is a
2 gateway node.

1 9. The network of claim 8 wherein said gateway node is a master network information server
2 node.

1 10. The network of claim 6 wherein the master network information server node is elected by
2 an election process comprising the steps of:

3 broadcasting an election datagram to said at least two information server nodes, said
4 election datagram comprising election criteria; and
5 electing a master network information server node in response to said election criteria.

1 11. The network of claim 6 wherein said master network information server node broadcasts a
2 declare datagram to detect multiple master network information server nodes using the same
3 transport protocol.

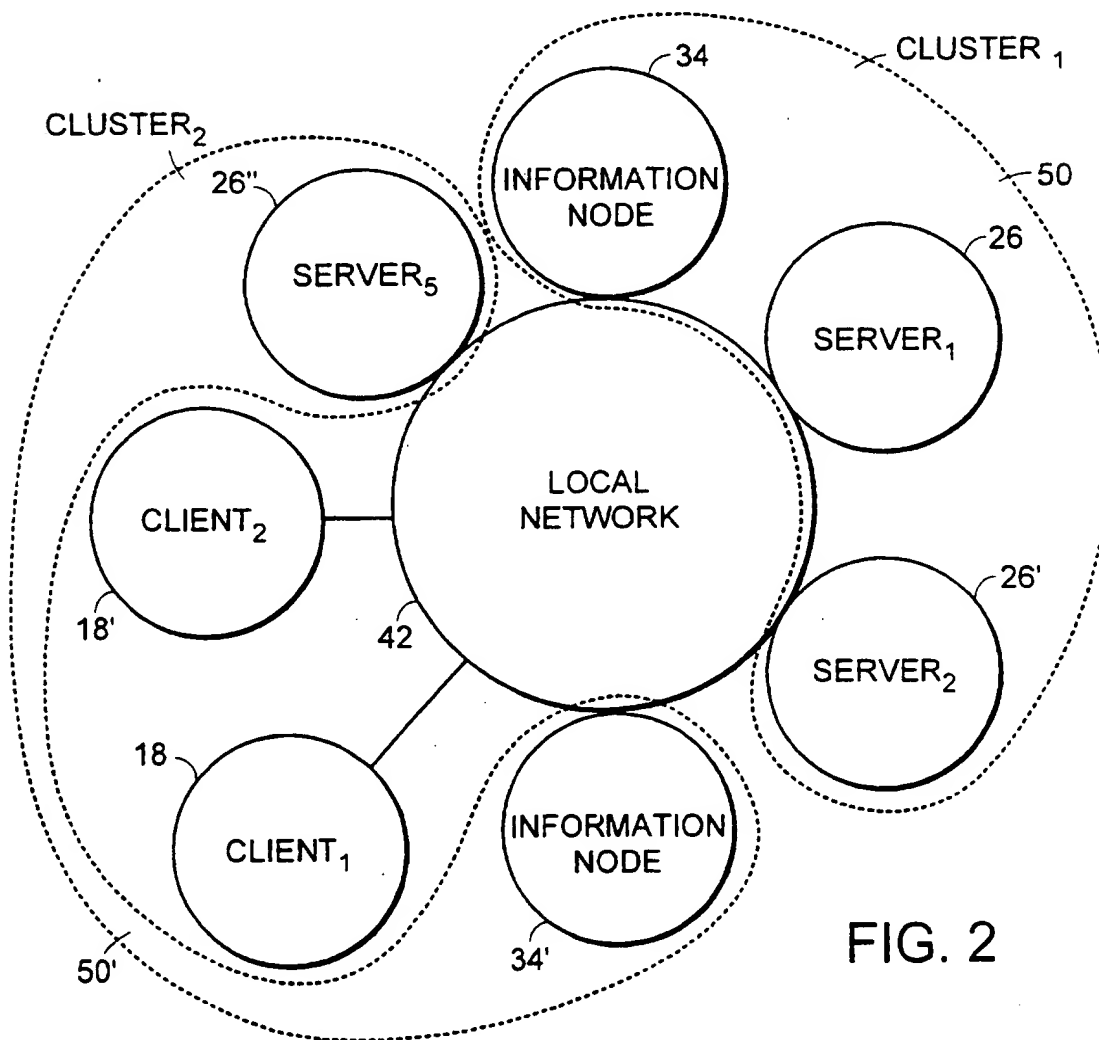

1 12. The network of claim 4 comprising a second information server node having a different
2 network transport protocol from said first information server node, wherein each of said
3 information server nodes is a master network information server node.

1 13. The election process of claim 10 wherein said election process occurs after an event
2 selected from the group of events consisting of a system reboot, a master network information
3 server node failing to respond with an acknowledge after being sent a server update datagram, a
4 master network information server node failing to respond to a client request, at least two master
5 network information server nodes configured with the same transport being detected, and a new
6 server node appearing on said network.

- 16 -

- 1 14. A method of connecting a client process to a server process on one of at least two
2 available servers, said method comprising the steps of:
- 3 requesting an available server from a network information server node; and
4 connecting to said server process on one of said at least two available servers in response
5 to said load levels as stored on said network information server node.

1/5



2/5




FIG. 3

3/5

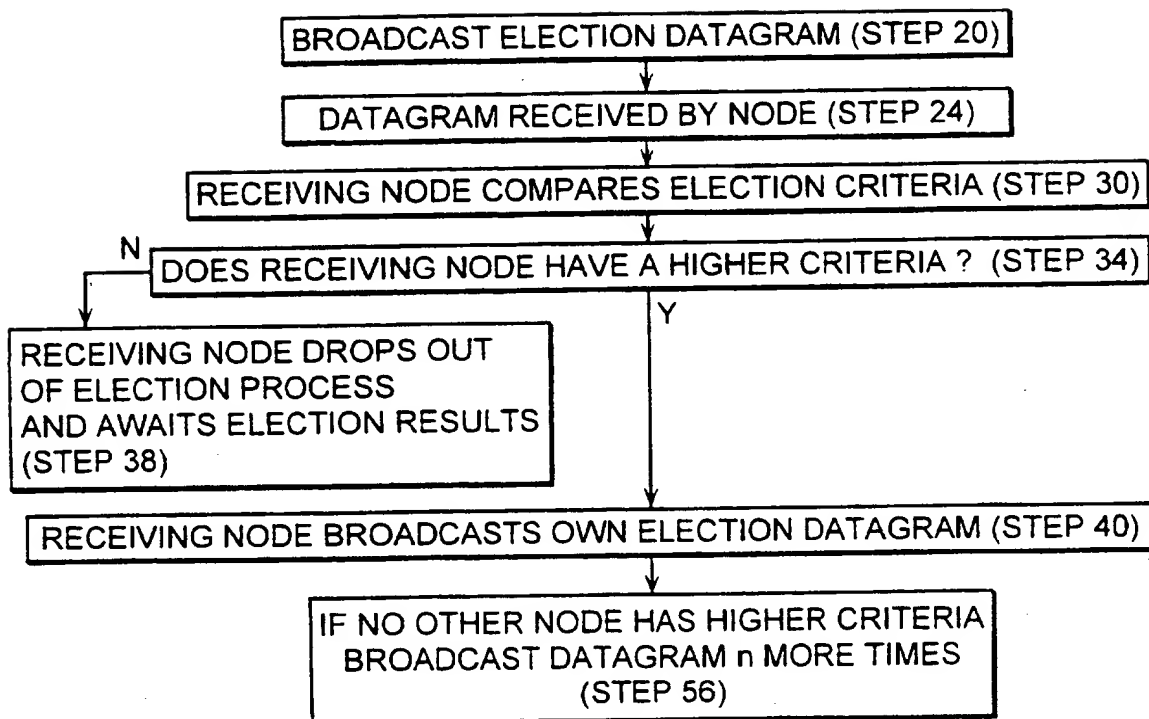


FIG. 4

4/5

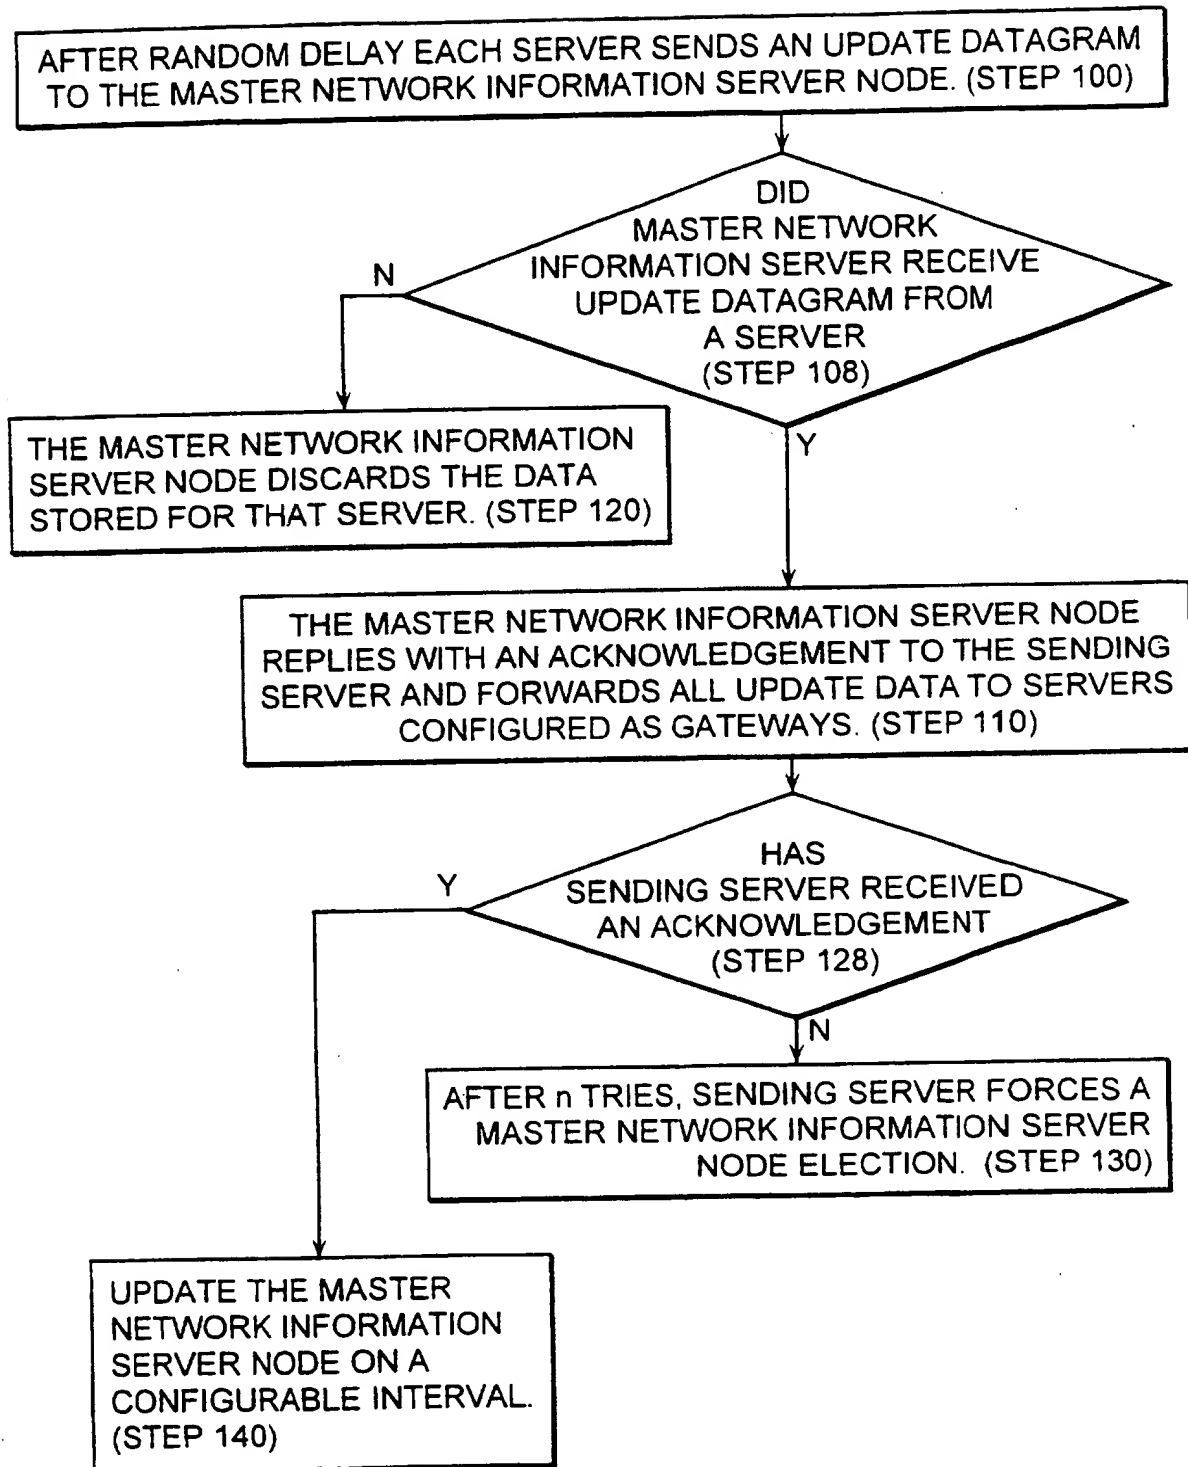


FIG. 5

5/5

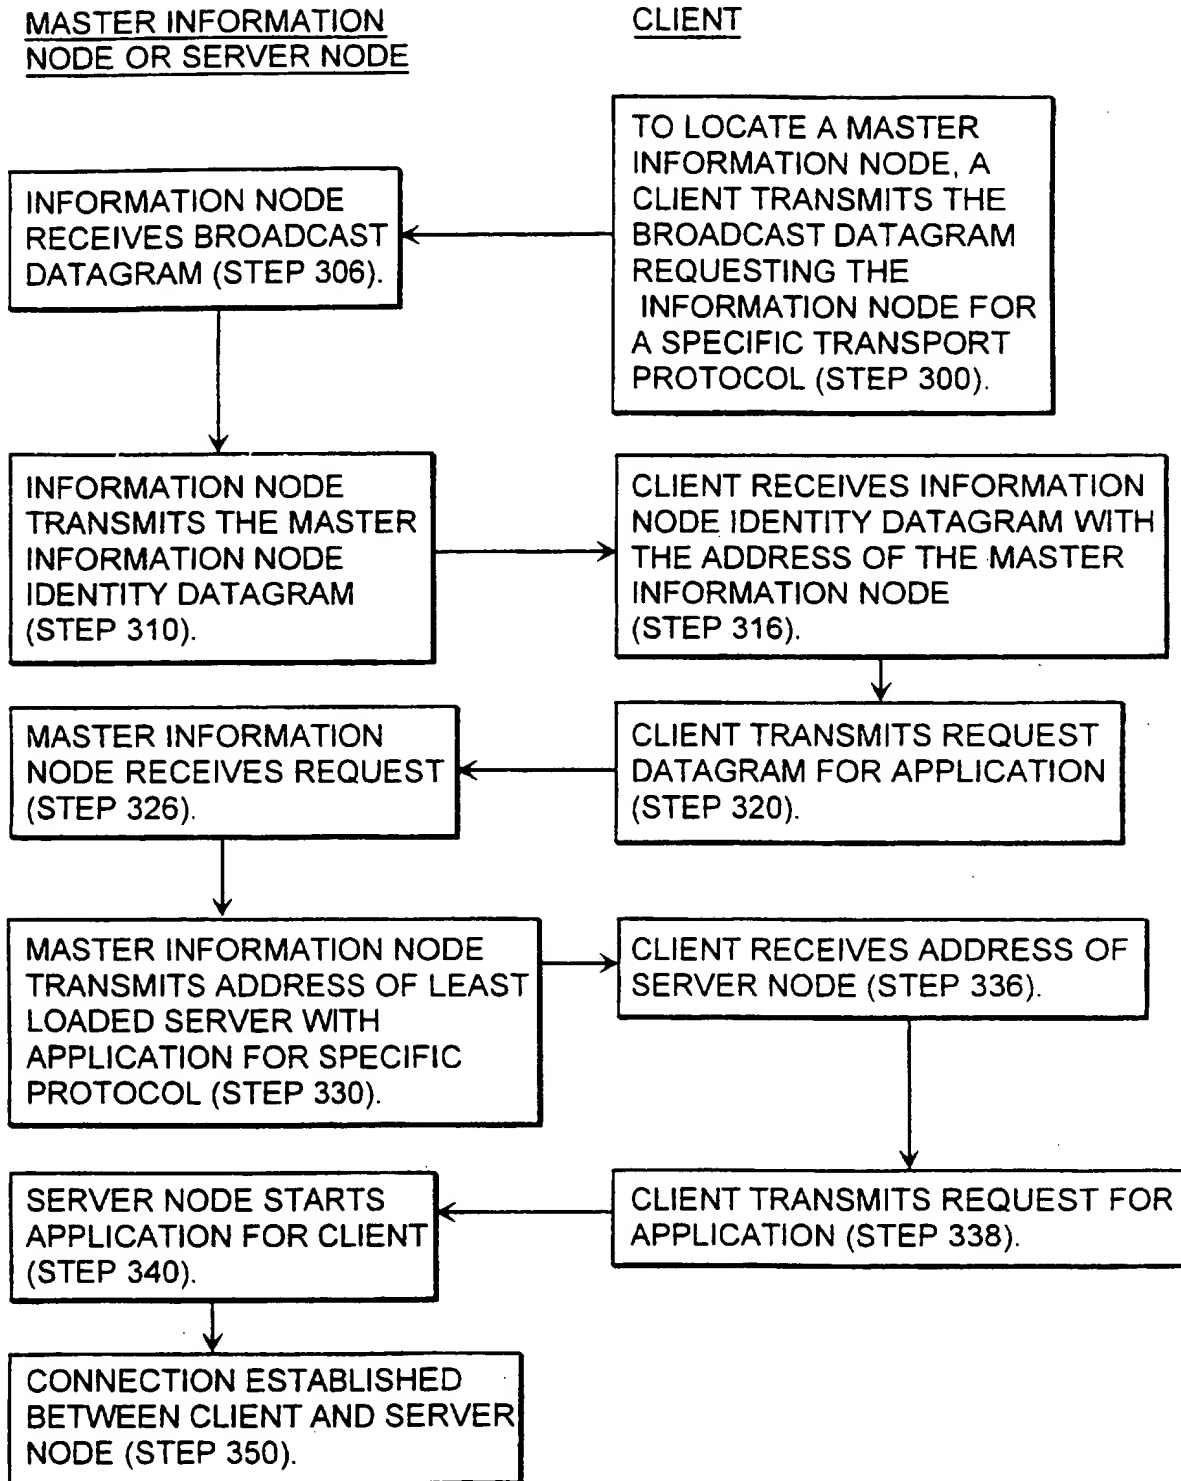
CLIENT REQUESTS

FIG. 6

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 97/01474

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06F9/46

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 648 038 A (IBM) 12 April 1995 see column 1, line 41 - column 2, line 6 see column 2, line 56 - column 3, line 54 see column 5, line 5 - line 39 ---	1-5,7,14
Y	EP 0 384 339 A (DIGITAL EQUIPMENT CORP) 29 August 1990 see column 2, line 50 - column 3, line 48 ---	1-5,14
Y	PATENT ABSTRACTS OF JAPAN vol. 095, no. 003, 28 April 1995 & JP 06 332782 A (HITACHI LTD;OTHERS: 01), 2 December 1994, see abstract --- -/--	1-5,14

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

Z document member of the same patent family

Date of the actual completion of the international search

5 June 1997

Date of mailing of the international search report

17. 06.97

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+ 31-70) 340-2040, Tx. 31 651 epo nl,
Fax (+ 31-70) 340-3016

Authorized officer

Michel, T

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 97/01474

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	IBM TECHNICAL DISCLOSURE BULLETIN, vol. 36, no. 1, 1 January 1993, pages 33-36, XP000333761 "OS/2 EE DATABASE MANAGER SQLJRA REMOTE CONTROL" see the whole document ---	1,8,9
A	EP 0 540 151 A (IBM) 5 May 1993 see page 2, line 42 - line 47 ---	1,10,13
P,X	EP 0 715 257 A (BULL SA) 5 June 1996 see the whole document -----	1-6,10, 13,14

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 97/01474

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0648038 A	12-04-95	GB 2281793 A JP 2505116 B JP 7093238 A	15-03-95 05-06-96 07-04-95
EP 0384339 A	29-08-90	AT 151183 T AU 611605 B AU 4996190 A AU 630291 B AU 7603391 A CA 2010762 A DE 69030340 D JP 3116262 A US 5341477 A	15-04-97 13-06-91 13-09-90 22-10-92 15-08-91 24-08-90 07-05-97 17-05-91 23-08-94
EP 0540151 A	05-05-93	JP 5216845 A US 5539883 A	27-08-93 23-07-96
EP 0715257 A	05-06-96	FR 2727540 A WO 9617297 A	31-05-96 06-06-96

AMENDED CLAIMS

[received by the International Bureau on 18 August 1997 (18.08.97);
original claims 1-6, 8, 10, 12 and 14 amended; new claims 15-17 added; remaining claims unchanged (3 pages)]



1. In a networked system including network information server nodes, one of the network information server nodes dynamically electable as a master network information server node, a network information server node comprising:
 - a first process for allowing dynamic election of the network information server node as the master network information server node, the first process receiving transmitted election criteria and comparing the received election criteria with stored election criteria;
 - a second process for constructing a load table after the network information server node is elected as the master information network server node, the load table constructed by storing a plurality of network addresses and a plurality of load information corresponding to a respective one of said plurality of network addresses in a memory element; and
 - a third process receiving client requests and responding to said client requests based on the constructed load table.
2. In a networked system including network information server nodes, one of the network information server nodes dynamically electable as a master network information server node, a network information server node comprising:
 - a first process for allowing dynamic election of the network information server node as the master network information server node, the first process receiving transmitted election criteria and comparing the received election criteria with stored election criteria;
 - a second process for constructing a load table after the network information server mode is elected as the master network information server mode, the load table constructed by storing a plurality of network addresses and load information associated with each of said plurality of network addresses in a memory element; and
 - a third process receiving client requests and responding to said client requests based on the constructed load table.
3. In a networked system including network information server nodes, one of the network information server nodes dynamically electable as a master network information server node, a network information server node comprising:
 - a first process for allowing dynamic election of the network information server node as the master network information server node, the first process receiving transmitted election criteria and comparing the received election criteria with stored election criteria;
 - a second process for constructing a load table after the network information server node is elected as the master network information server node, the load table constructed by storing an entry corresponding to a respective node in a memory element, each entry including:

a network address of a respective one of said plurality of network nodes, and
load information associated with said respective one of said plurality of network nodes;
a third process receiving client requests and responding to each of said client requests with entries of said load table corresponding to network nodes which satisfy said each of said client requests.

4. A network comprising:
at least two server nodes;
a master network information server node in communication with said at least two server nodes, said master network information server node comprising a memory comprising a data table and at least one election criteria for allowing said network information server node to be dynamically selected, said data table having an entry for each of said at least two server nodes, each of said entries including: an address corresponding to a respective one of said at least two server nodes; and load information corresponding to said respective one of said at least two server nodes; and
a client node in communication with said master network information server node and said at least two server nodes, said client node executing an application on one of said at least two server nodes in response to said load information obtained from said master network information server node.
5. The network of claim 4 further comprising a plurality of client nodes in communication with said master network information server node.
6. The network of claim 4 further comprising at least two network information server nodes wherein one of said at least two network information server nodes is a master network information server node.
7. The network of claim 4 wherein said at least two server nodes are grouped into clusters.
8. The network of claim 6 wherein one of said at least two network information server nodes is a gateway node.
9. The network of claim 8 wherein said gateway node is a master network information server node.
10. The network of claim 6 wherein said master network information server node is elected by an election process comprising the steps of:
broadcasting an election datagram to said at least two network information server nodes, said election datagram comprising election criteria; and
electing a master network information server node in response to said election criteria.

11. The network of claim 6 wherein said master network information server node broadcasts a declare datagram to detect multiple master network information server nodes using the same transport protocol.
12. The network of claim 4 comprising a second network information server node having a different network transport protocol from said first network information server node, wherein each of said network information server nodes is a master network information server node.
13. The election process of claim 10 wherein said election process occurs after an event selected from the group of events consisting of a system reboot, a master network information server node failing to respond with an acknowledge after being sent a server update datagram, a master network information server node failing to respond to a client request, at least two master network information server nodes configured with the same transport being detected, and a new server node appearing on said network.
14. A method of connecting a client process to a server process on one of at least two available servers, said method comprising the steps of:
 - determining which of the at least two server nodes is the master network information node;
 - requesting an available server from the master network information server node; and
 - connecting to said server process on one of said at least two available servers in response to said load levels as stored on the master network information server node.
15. The network information server node of claim 1 wherein said first process, said second process, and said third process comprise a single process.
16. The network information server node of claim 2 wherein said first process, said second process, and said third process comprise a single process.
17. The network information server node of claim 3 wherein said first process, said second process, and said third process comprise a single process.

412JDL1545/13.393001-1



2/5




FIG. 3

3/5

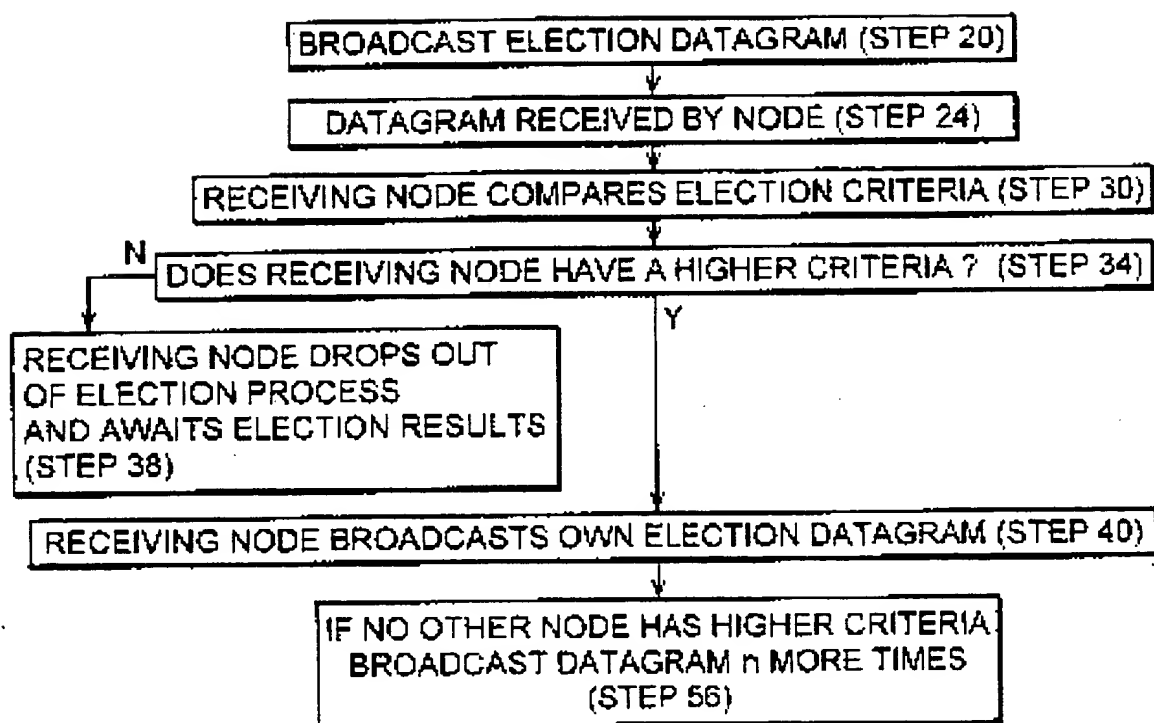


FIG. 4

4/5

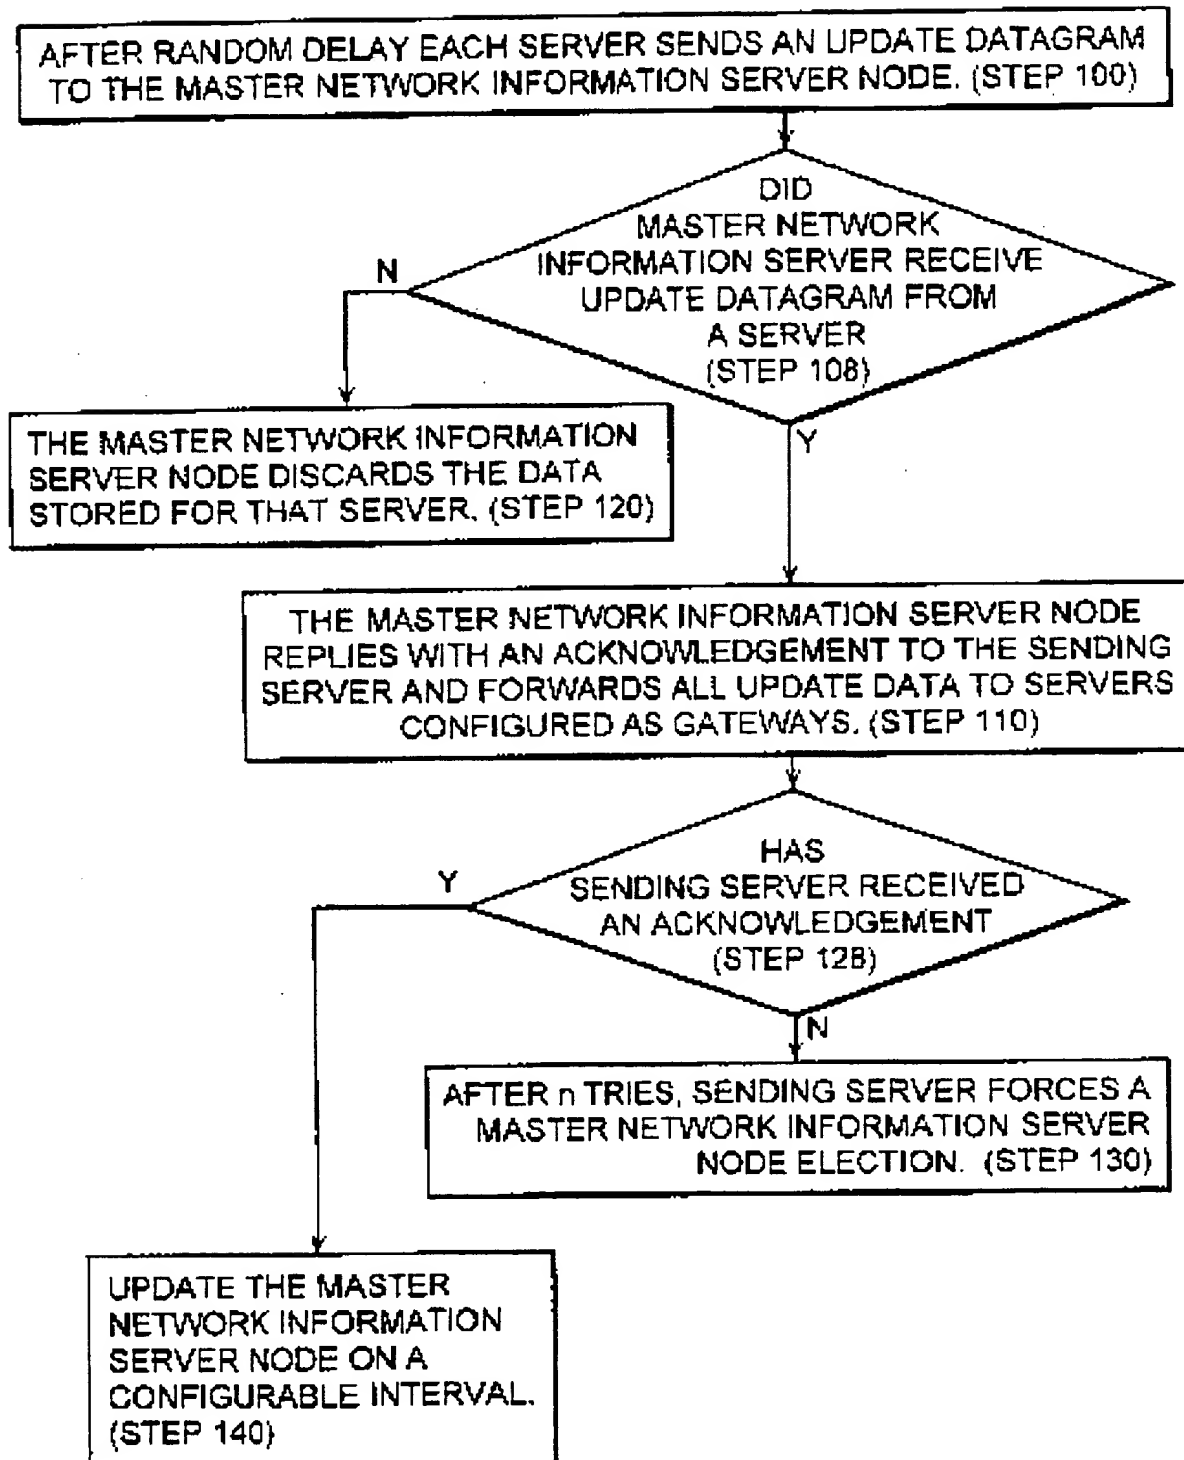


FIG. 5

5/5

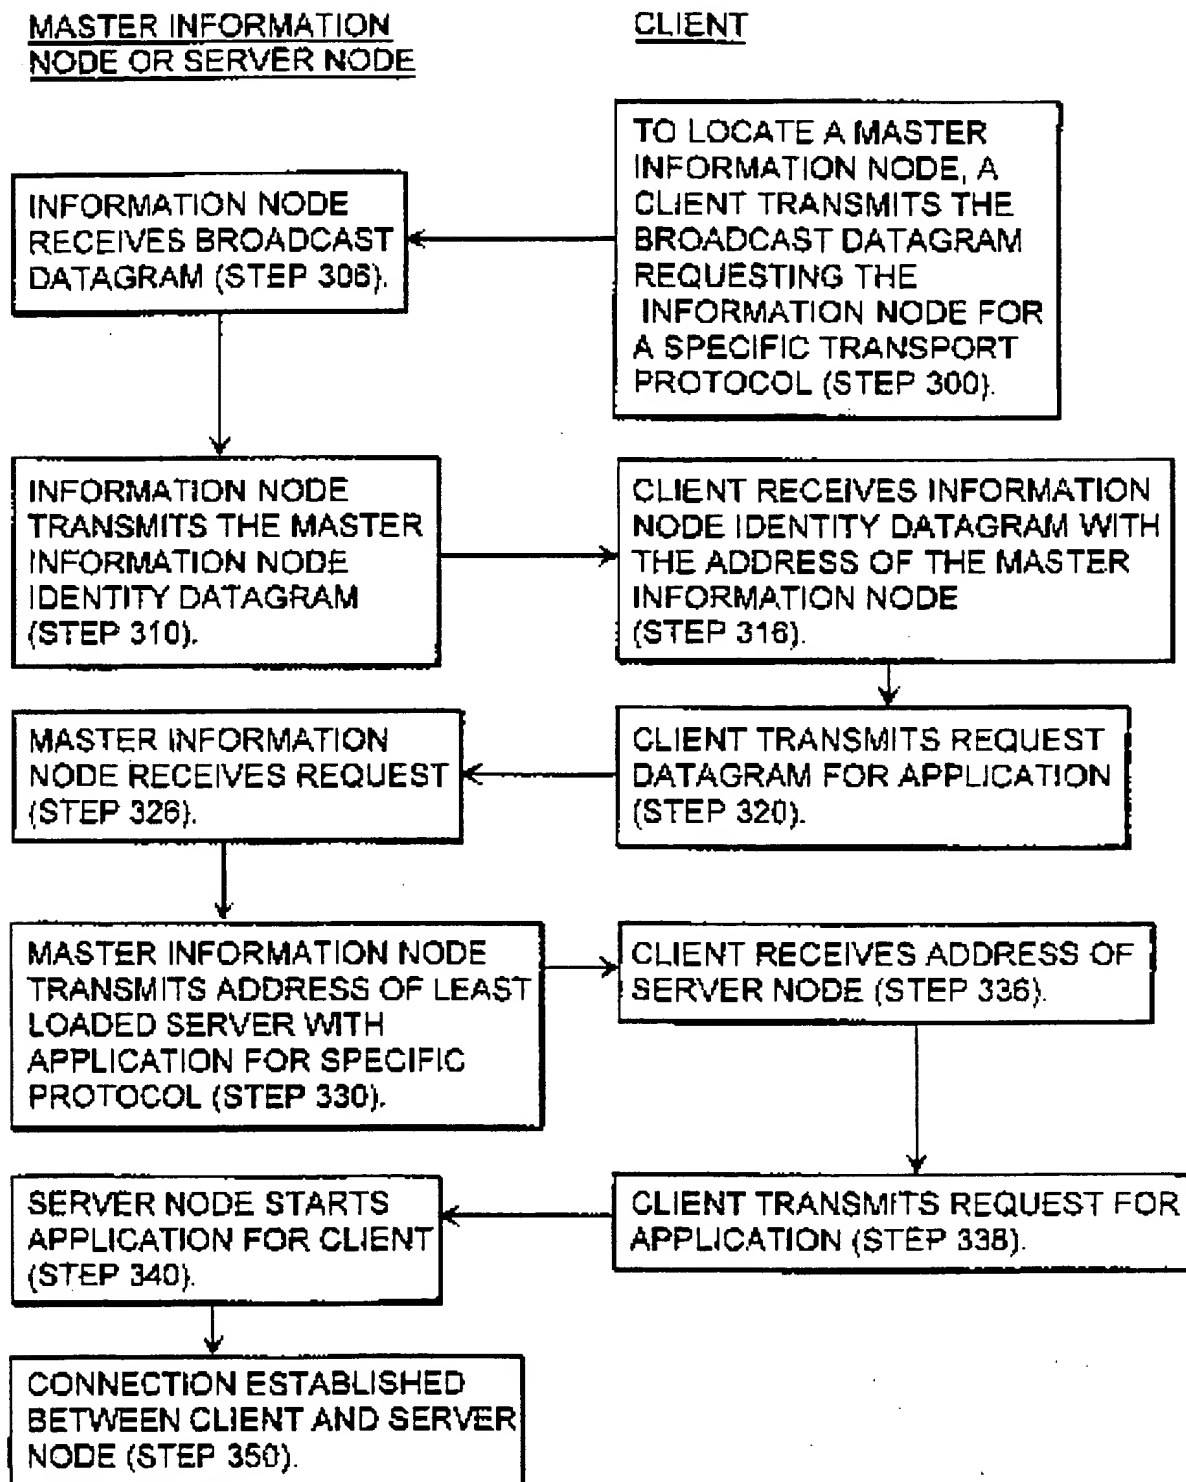
CLIENT REQUESTS

FIG. 6